

FOIVOS TSIMPOURLAS

foivos.co.uk

fivosts@gmail.com

San Francisco, CA, United States

PROFESSIONAL EXPERIENCE

OpenAI

Research Scientist

July 2023 - Present

San Francisco, CA

- I was part of the reasoning team until March 2025 and was one of the foundational contributors to the releases of Deep Research, O1 and O3.
- I am part of the Health AI team and contributed to releasing HealthBench and ChatGPT Health.

TikTok - ByteDance

Research Scientist, Contractor

Oct 2022 - Jan 2023

London, UK

- I developed and published the first LLM-based tool for fault localization in Go concurrent test executions.

Facebook AI Research

Research Scientist Intern

May 2022 - August 2022

Menlo Park, CA

- I created the first directed code generator using Deep RL, LLMs and sampling guidance towards compiler features.
- I published my work in PACT and integrated it into CompilerGym, a state-of-the-art research tool for RL.

Athens Digital Arts Festival

Research Scientist, Contractor

Oct 2020 - Dec 2021

Remote

- I secured the research grant and led the research effort for "AITrace", a real time recommender for museums.
- I developed the AI for the project's pilot and reported our experimental findings to the Greek Deputy of Culture.

ARM Ltd.

Software Engineer, Intern

Sept 2019 - Jan 2020

Cambridge, UK

- I was the DRI for the first distributed GICv4 interrupt controller's software model for ARM-based clusters.
- I collaborated with Architecture researchers and Modeling engineers to implement and ship to production.

EDUCATION

PhD in Informatics

University of Edinburgh (UoE)

2019 - 2023

Thesis: "Deep Language Models for Software Testing and Optimisation"

Supervisors: Ajitha Rajan, Hugh Leather, Pavlos Petoumenos

- My research contributed to the integration of AI into compilers using LLMs, RL and the LLVM framework.

MSc by Research, CDT in Pervasive Parallelism

University of Edinburgh (UoE)

2018 - 2019

Distinction

Thesis: "Learning to Encode and Classify Test Executions"

Supervisors: Ajitha Rajan, Christophe Dubach

- My research focused on deep learning for compile runtime analysis and program correctness prediction.

MEng, Electrical Engineering & Computer Science

National Technical University of Athens

2011 - 2018

Grade: 8.1/10

Thesis: "Resource Management Techniques for Embedded Architectures executing Deep Neural Networks"

Supervisors: Dimitrios Soudris

- I improved GoogleNet's execution time by 16% on Intel's Myriad2 with a state of the art hardware acceleration algorithm.

PUBLICATIONS

gpt-oss-120b & gpt-oss-20b model card <i>OpenAI and F. Tsimpourlas</i>	2025 <i>Arxiv</i>
Healthbench: Evaluating large language models towards improved human health <i>RK Arora, J. Wei, RS Hicks, J. Quiñonero-Candela, F. Tsimpourlas, M. Sherman, M. Shah, A. Vallone, A. Beutel, J. Heidecke, K. Singhal</i>	2025 <i>Arxiv</i>
Competitive Programming with Large Reasoning Models <i>OpenAI and F. Tsimpourlas</i>	2025 <i>Arxiv</i>
Openai o1 system card <i>OpenAI and F. Tsimpourlas</i>	2025 <i>Arxiv</i>
Go-Oracle: Automated Test Oracle for Go Concurrency Bugs <i>F. Tsimpourlas, C. Rosuero, C. Peng, P. Yang, A. Rajan</i>	2024 <i>ICSE</i>
Deep language models for software testing and optimisation <i>F. Tsimpourlas</i>	2023 <i>PhD Thesis</i>
BenchDirect: A Directed Language Model for Compiler Benchmarks <i>F. Tsimpourlas, P. Petoumenos, M. Xu, C. Cummins, K. Hazelwood, A. Rajan, H. Leather</i>	2022 <i>Arxiv</i>
BenchPress: A Deep Active Benchmark Generator <i>F. Tsimpourlas, P. Petoumenos, M. Xu, C. Cummins, K. Hazelwood, A. Rajan, H. Leather</i>	2022 <i>PACT</i>
Embedding and Classifying Test Execution Traces using Neural Networks <i>F. Tsimpourlas, G. Rooijackers, A. Rajan, M. Allamanis</i>	2021 <i>IET Software</i>
Supervised Learning over Test Executions as a Test Oracle <i>F. Tsimpourlas, A. Rajan, M. Allamanis</i>	2021 <i>SACSE</i>
AITrace: A Deep Learning Recommender for Museum Visitors <i>F. Tsimpourlas, F. Salmouka, M. Markellou</i>	2021 <i>HFRI Greece</i>
A Design Space Exploration Framework for Convolutional Neural Networks Implemented on Edge Devices <i>F. Tsimpourlas, L. Papadopoulos, A. Bartsokas, D. Soudris</i>	2018 <i>CODES+ISSS</i>

TECHNICAL SKILLS

Programming Languages	C++, C, Python, Bash, Assembly (x86, ARM, SHAVE-VLIW)
Machine Learning	Pytorch, TensorFlow, WandB, brix, Snowflake
Hardware Description	SystemC, LISA+
Compilers	LLVM, Libtooling, GCC, Clang
Tools	Valgrind, GDB, CMake, Bazel

ACADEMIC ACTIVITIES

Peer Reviews

- **FSE** - Reviewed ML on compiler engineering papers on main conference research track. 2022
- **ASE** - Reviewed ML to software testing papers on main conference research track. 2019
- **PATMOS** - Reviewed hardware acceleration papers for ML. 2018

Teaching

- Tutor, **Software Testing**, University of Edinburgh. Voted best tutor by students twice. 2019-2022
- Tutor & Marker, **Microprocessors Systems Lab**, NTUA 2017 - 2018
- Teaching Assistant, **Embedded Systems Design**, NTUA 2017 - 2018

Talks

- **PACT** - Delivered conference presentation for accepted publication as 1st author. 2022
- **HFRI Greece** - As AITrace's research lead, I presented our research roadmap and findings to HFRI's research committee including the Greek Deputy of Culture. 2021
- **SACSE** - Delivered conference presentation for accepted publication as 1st author. 2021

SCHOLARSHIPS AND AWARDS

· **HFRI grant no. 00782 for “AITrace” research project.** *2021*

Our research proposal secured €150,000 of research funding. 100 awarded research teams among thousand applicants.

· **PhD scholarship, EPSRC grant EP/L01503X/1.** *2018*

I was 1 of 8 students accepted to “CDT in Pervasive Parallelism” from thousands of applicants.

· **High school award for Panhellenic university entrance exams score.** *2011*

I achieved an average score of 97.48% placing in the top 1% among 90,000 students in the Panhellenic university entrance exams. I was one of very few students nationwide that scored 100% in all 5 science subjects.